The International Center for Research on the Management of Technology

The Emergence of a New Supercomputer Architecture

Alan N. Afuah
James M. Utterback

July 1990 3215-90-BPS

Sloan WP# 90-3215BPS

Sloan School of Management
Massachusetts Institute of Technology
38 Memorial Drive, E56-303
Cambridge, MA 02139

Abstract

A model on the dynamics of innovation among multiple productive units, and other innovation models are used to examine the emergence of a new supercomputer architecture. Data on entry and exit of firms producing supercomputers having three distinctive architectures appear to conform to the main hypotheses of the models examined. Based on these data, the authors speculate that the new generation of massively parallel supercomputers will replace the currently accepted von Neumann architecture for supercomputers though the ascendant dominant architecture cannot yet be spelled out in detail. Complementary assets, especially software, and chance events will all play a part in determining which massively parallel design may ultimately be used for most supercomputing applications.

1 Introduction

Over the years, computer-intensive applications like aeronautical structural analysis, finite element analysis, atomic energy and nuclear fusion, and numerical weather forecasting have used supercomputers to provide the computation power needed. Until the 1980s, supercomputers were designed using the classic sequential computer architecture of John von Neumann, with incremental innovations like pipelining and vector processors added to boost speed. Then innovators decided to attack the supercomputer market with new designs: minisupercomputers in the early 1980's, and massively parallel computers -- a radical innovation -- in the mid-80s.

This paper examines the development of traditional supercomputers and the present competition between them and massively parallel computers

(MPC) using the Abernathy & Utterback [1978] dynamic model of innovation, and the Utterback & Kim [1986] hypotheses on discontinuous change in a product; and concludes that MPCs are strongly invading the traditional supercomputers. It then uses the Utterback [1987] model on the dynamics of innovation among multiple productive units, to examine the different productive units -- traditional supercomputers, minisupercomputers and massively parallel computers. The analysis concludes that dominant designs are about to emerge and will be of the MPC architecture; but cannot tell exactly what particular designs of the MPC architecture will dominate.

2 Three Models of Innovation

In their dynamic model of innovation, Abernathy and Utterback [1] describe the evolution of products and processes from a *fluid state,* through a *transitional* state to a *specific* state. In the *fluid* state, the performance requirements for new products and market needs are not well defined and that means that the source of innovation or technology is often users [von Hippel 1988]. The fluid state is a state of flux with a rapid pace of technological innovation. The basis of competition is on performance and technological characteristics. As time goes on and some standardization takes place, a dominant design emerges and production volumes increase, with design goals well-articulated, leading to the *specific* state. The product is now a commodity product with the basis of competition shifting from performance to price and cost considerations. Innovation is less rapid and more incremental in nature.

In the specific state, the organizational structure and control are more formal with technology planning and forecasting being formally delegated tasks. The model applies to assembled products; although lately, Utterback and Nolet [1987] have done some work on product and process change in non-assembled products like flat glass, rayon and electric power generation. The model implies that effective management of large companies means management of a portfolio of businesses from all three states -- fluid, transitional and specific. Details of the characteristics of each state are shown in Table 1

TABLE 1

The Abernathy and Utterback Dynamic Model of innovation.

	FLUID PATTERN	TRANSITIONAL PATTERN	SPECIFIC PATTERN
COMPETITIVE EMPHASIS ON	Functional product performance	Product variation	Cost reduction
INNOVATION STIMULATED BY	Information on users' needs and users' technical inputs	Opportunities created by expanding internal technical capability	Pressure to reduce cost and improve quality
PREDOMINANT TYPE OF INNOVATION	Frequent major changes in products	Major process changes required by rising volume	Incremental for product and process, with cumulative improvement in productivity and quality
PRODUCT LINE	Diverse, often including custom designs	Includes at least one product design stable enough to have significant production volume	Mostly undifferentiated standard products
PRODUCTION PROCESS	Flexible and inefficient; major changes easily accommodated	Becoming more rigid, with changes occurring in major steps	Efficient, capital intensive and rigid; cost of change is high.
EQUIPMENT	General purpose requiring highly skilled labor	Some sub-processes automated creating "islands of automation"	Special-purpose mostly automatic with labor tasks mainly monitoring and control
MATERIALS	Inputs are limited to generally available materials	Specialized materials may be demanded from some suppliers	Specialized materials will be demanded, if not available, vertical integration will be extensive
PLANT	Small scale, located near user or source of technology	General-purpose with specialized sections	Large-scale, highly specific to particular products
ORGANIZA-TIONAL CONTROL	Informal and entrepreneurial	Through liaison relationships, project and task groups.	Through emphasis on structure, goals, and rules

Reprinted with permission from, Abernathy, William J and James M. Utterback, "Patterns of innovation," *Technology Review*, Massachusetts Institute of Technology, vol. 80, No. 7, June/July 1978.

While the dynamic model of innovation just discussed is about evolutionary change, the Utterback & Kim [1986] model is about radical change – change that renders a company's investments in design know-how, manufacturing skills, plant and equipment, and other technical skills and knowledge useless. In particular, their model focuses on two productive units: one from an older technology that is likely in the *specific* state discussed above, and another from a newer technology, also very likely to be in the *fluid* state; with the newer technology invading the old.

The invading technology often has at least one performance charactistic that is superior to (or potentially so) that of the invaded technology. The invading companies will go after the niche markets for which their technology ofers of the performance advantage. Often, the invading company views the new technology just as a replacement for the old and does not see the potential additional uses. A classic example is invasion of the vacuum tube market by transistors.

Tushman and Anderson [1986] have also studied the phenomenon of radical innovation, classifying innovation as either competence-enhancing or competence-destroying for vacuum tube makers, but competence enhancing for computer makers. The Utterback & Kim hypothesis concerning discontinuous change in a product are shown in Table 2.

TABLE 2

The Utterback & Kim Hypotheses Concerning Discontinuous Change in a Product (James M. Utterback and Linsu Kim, "Invasion of a stable business by by radical innovation." *The Management of Productivity and Technology in Manufacturing*, by Kleindorfer, P. R.)

Invaded unit will	Invading unit will
Be an established competitor in the product market segment in question, often producing a standard product in large volume.	Tend to be a new entrant, either a new enterprise or a larger firm carrying its technology into a new market area.
Generally exhibit the specific or stable pattern of characteristics of the Abernathy and Utterback model referenced earlier.	Generally exhibit the pattern of fluid or flexible characteristics of the Abernathy and Utterback model referenced earlier.
Strongly perceive the continuing advantages of its product in terms of both performance and cost in contrast to the disadvantages of the invading innovation.	Stress the unique advantages of its innovation for some demanding applications or market niches, often not addressed by the established technology.
Tend to view the invading unit as a substitute.	Tend to view its innovation as expanding the existing market or broadening the existing range of product uses. It may later become a substitute.
Often experience a continuing expansion of demand for its product over a considerable time (a decade or more).	May, by introducing its innovation actually strengthen demand for the established technology as well for a time through various complementaries.
Tend to be limited in its responses by large investments in people, equipment, plant, materials and technical knowledge which are irrelevant for success in the invading technology.	Tend to have established strengths in the requisite technology or to acquire them via rapid expansion.
When clearly threatened by the invading innovation, invaded unit will redouble its investments in an effort to improve its traditional technology. This will result in a period of renewed search and creativity and often in dramatic improvement of what had been a stable technology.	Often be in a superior competitive position with respect to skills in the new technology, production and market understanding, and rates of improvement by the time the established unit begins to respond.
Fail to survive the invasion in most cases However, an established unit which has diverse and significant technical skills. and which dominates its market, may be able to successfully acquire the invading technology.	Ultimately dominate the invaded product market segment. However, it may fail if it does not appropriately evolve beyond its initial innovation, and it may be acquired by the established unit, or, in cases, even overwhelmed by creative response of the old technology to the invasion.
Continue to be produced for some specialized market niches for which the invading technology is less viable	In turn, be invaded by yet another wave of radical innovation.

The dynamics of innovation among multiple productive units:

Abernathy & Utterback [1978], and Utterback & Kim [1986] respectively discussed hypotheses about the evolution of a productive unit and the discontinuities from one productive unit to another following a radical innovation. The hypotheses assume that at any one time, only one productive unit exists in a particular market segment. But in the real world, this assumption does not hold. The Utterback [1987] model on the dynamics of innovation on multiple productive units relaxes the restrictive assumption and examines the implications of these hypotheses on a set of productive units competing in the same market facing similar constraints.

The model contends that the rate at which firms enter or exit an industry parallels the product innovation in that industry. Thus, in the fluid state where product requirements are still ambiguous, there is expected to be rapid entry of firms with very few or no failures. As the industry enters the transitional state, product requirements become more universal as a dominant design emerges. Fewer firms enter and many exit or merge. In the specific state, there emerges an oligopoly controlling a fairly consistent share of the market. Mueller and Tilton (1969) present similar arguments.

According to Utterback (1987) , creative synthesis of a new product innovation by one or a few firms may result in a temporary monopoly, high unit profit margins and prices, and sales of the innovation in those few market niches where it possesses the greatest performance advantage over competing products. As volume of production and demand grows, and as a wider variety of applications is opened for the innovation, many new firms enter the market with similar products.

The appearance of a dominant design shifts the competitive emphasis to favor those firms with a greater skill in process innovation and process integration, and with more highly developed internal technical and engineering skills. Many firms will be unable to compete effectively and will fail. Others may possess special resources and thus merge successfully with the ultimately dominant firms, whereas weaker firms may merge and still fail.

Eventually, the market reaches a point of stability, corresponding to the specific state, in which there are only a few firms -- four or five is a typical number from the evidence reviewed by Utterback [1987] -- having standardized or slightly differentiated products, and stable sales and market shares. A few small firms may remain in the industry, serving specialized market segments, but, as opposed to the small firms entering special segments early in the industry, they have very little growth potential. Thus, it is important to distinguish between small surviving firms and small firms that are new entrants, and to keep in mind that the term "new entrants" includes existing larger firms moving from their established market or technological base into a new product area.

The development of a set of productive units is expected to begin with a wave of entry gradually reaching a peak about when the dominant design of the major product emerges, and then rapidly tapering off. This sequence is followed by a corresponding wave of firms exiting from the industry. The sum of the two waves -- entries and exits -- will yield the total number of participants in the product market segment at any time. Therefore the number of participants in an industry can be represented by a curve that starts with a gentle rise representing the first few fluid productive units entering the business followed by a much sharper rise that represents a wave of imitating firms. The point at which a dominant design is introduced in the industry is followed by a sharp decline in the total number of participants until the curve of total participants reaches the stable condition with a few firms sharing the market.

The model implies that it is theoretically possible to enter a market segment at any of theses stages by formulation and execution of the right strategy: in the fluid state, stress a high degree of product innovation; in the transitional state, stress process innovation and process integration; and in the specific state by having the financial strength to invest in a plant with the most scale economies and at the best location. However, no U.S. data set examined to date contains an example of a successful surviving firm which entered after introduction of a dominant design. Utterback (1987) speculates that compounding evidence will support the observation that emergence of a dominant design will result in the beginning of a wave of exits from an industry segment, and thus will correspond with a peak in the number of firms participating in that segment.

3 The Supercomputer Industry Examined

The Technology and Products

Supercomputers are generally described as the most powerful computational systems available at any given time [Supercomputing 1987] with speeds today in the 100 MFLOPS (Million FLoating Point Operations per Second) and higher range. Three major design technologies have been used to build supercomputers: Sequential, vector, and parallel processing.

Sequential computers have only one CPU (Central Processing Unit, the "engine" or brain of the computer) that can only do one thing at a time. As such, even problems that are inherently parallel like matrix multiplication have to be broken down serially in order to be processed by such machines.

Vector processors allow simultaneous computation for some problems, and can thus be used to speed up programs that have a vector- or matrix-like structure.

In parallel processing, many processors -- hundreds or thousands in the case of massively parallel computers -- are put on one job to get the job done faster than one processor or a few processors -- the structure of the job permitting. Massively parallel processing is radically different from vector processing in terms of both programming and potential applications.

Supercomputer makers

Makers of supercomputers can be grouped into three categories of competitors: i) the makers of traditional von Neumann[1] architecture machines with associated incremental innovations of vector processors and 1-8 processors in parallel, ii) Minisupercomputer makers and iii) massively parallel computer makers. Products from categories *i* and *ii* use some combination of the sequential and vector processor technologies, and are to a large extent, still of the von Neumann architecture, while products from *iii* are fundamentally parallel and conceptually, radically different from the von Neumann architecture.

Traditional Supercomputer makers include Cray Research Inc., Fujitsu, Hitachi, NEC, IBM, Supercomputer Systems Inc. (SSI), CDC, and the now defunct ETA Systems, and Denelco; with the market being dominated by Cray Research Inc.. Their machines are predominantly of the von Neumann sequential architecture with vector processors to boost performance. They use the fastest chips available and have elaborate liquid-cooling systems. For most of the literature on high performance computing, supercomputers consist of only machines from this category.

Minisupercomputers utilize the von Neumann sequential architecture with the associated incremental innovations of pipelining, and vector

[1]The architecture used in most of today's computers is often attributed to John von Neuamann's mid-1940s architecture. In that architecture, the Central Processing Unit (CPU) of the computer fetches an instruction (data) from a central store (main memory), operates on it (for example, add or subtract), and returns the results into the main memory. Only one CPU is used, and that one CPU can do only one thing at a time.

processing of traditional supercomputers, but with some important differences: Minisupercomputers are cheaper, provide 25 to 35 percent [Kelley 1988] the performance of Cray-type traditional supercomputers, offer lower price for the performance provided and lend themselves to those low-end applications that don't need the higher performance of a Cray, let alone its price. They use CMOS (Complementary Metal Oxide Semiconductor) chips that are less expensive and consume less power than the power-demanding but faster ECL [2] chips used in traditional designs. This results in cheaper systems that are air-cooled.

Massively Parallel Computers (MPC) are architecturally and conceptually very different from the other two categories of supercomputers. They use hundreds and sometimes thousands of processors, each of which doesn't have to be very fast. These machines use slower, proven chips and other components in their designs, and are air-cooled. See Afuah [1990] for a more detailed description of the technology and products in the supercomputer industry.

Data

The primary sources for the data were product managers, product analysts, and engineers (design, applications, and marketing) in the supercomputer industry. Whenever possible, employees who had worked both at a traditional supercomputer company (incumbent) and then at a massively parallel computer company or minisupercomputer company were used as sources. Supercomputer customers and suppliers also proved to be

---

[2] ECL stands for Emitter Coupled Logic, a silicon chip technology that results in the fastest silicon chips, but which consume a lot of power and are very expensive compared to chips from other silicon technologies.

very useful sources. Secondary data sources – mostly trade and scientific journals and magazines – were also used.

Entries and exits into the supercomputer industry were noted as follows: An entry occurred in the year that the company announced its entry. (Shipment of the first product usually followed zero to three years later). The exit year is the year the company announced it was getting out of the business and stopped production. Acquisition of a company by another company that already made supercomputers was counted as one exit. The merger of two companies to form one supercomputer company was also counted as one exit and no entry. Table 3 summarises all the exits and entries.

TABLE 3
Supercomputer makers -- Entries and Exits

Traditional Supercomputer makers	Entry	Exit date	Comment or Max. speed in 1990
IBM	1955	1971	
UNIVAC	1955	1958	
CDC	1962	1983	
Burroughs	1971	1974	
Texas Inst.	1972	1975	
Cray Research	1972		2.7 GFLOP
ETA Systems	1983	1989	
Denelcor	1982	1985	
Fujitsu	1982		4 GFLOP
Hitachi	1982		3 GFLOP
NEC (HNSX)	1983		22 GFLOP
IBM	1986		1.6 GFLOP
SSI	1988		N/A
Cray Computer Corp.	1989		16 GFLOP (in 1992)
Minisupercomputer Makers			
Floating Point Systems	1971		
Ardent	1985	1989	Merged with Stellar
Stellar	1985	1989	Merged with Ardent
Scientific Comp. Systems	1983	1989	
Convex Computing Corp.	1982		
Alliant Computer Systems	1982		
Multiflow Computer	1984	1990	
Cydrone	1984		
Encore Computers (Hydra)	1985		
Inter'l Parallel Computers	1980		
ELXSI	1982		
Stardent	1989		Merger of Stellar and Ardent
Supertek	1986	1990	Bought by Cray Research
Massively Parallel Computer makers			
Thinking Machines	1985		
Goodyear Aerospace	1983		
Bolt Beranek and Newmann	1984		
Intel Scientific Computers	1985		
N-cube	1983		
Floating Point Systems	1986	1990	But still in minisupers
MasPar	1988		
Meiko	1985		
Evans & Suntherland	1985	1989	
Active Memory Technology	1986		
Ametek	1983	1990	
Myrias	1983		

4 Application of the models to Supercomputers

The Abernathy & Utterback dynamic model of innovation can be used to establish that traditional supercomputer designs have the characteristics of the specific state, and massively parallel computers (MPC) have those of the fluid state. The patterns and hypotheses of discontinuous change developed by Utterback & Kim are then examined to show that MPCs also display the other characteristics hypothesized for an invading technology. Finally the Utterback Model on the dynamics of innovation among multiple productive units is used to conclude that the emerging dominant supercomputer designs will likely be of the MPC architecture.

There is nothing fundamentally different about the architecture of minisupercomputers (compared to the traditional designs) that can make them outperform the other two categories and become the dominant design. As such, the discussion will be limited to the MPCs and the Cray-type traditional supercomputers. Minisupercomputers remain an effective competitive weapon in the low-end segment of the market and not a replacement for MPCs or traditional supers. If anything, they too will be replaced by MPCs.

Traditional Supercomputers are in the Specific State of the Dynamic Model of Innovation:

For traditional supercomputers based on the von Neumann architecture, competitive emphasis has been on price/performance -- the more FLOPS or MIPS per dollar the better. To increase the performance of their machines, tradtional supercomputer makers have emphasized incremental innovations in the product, process, and materials. Still building around the von Neumann architecture, they have added such features as pipelining, vector processing and limited parallel processing (large-grained). NEC, for example, has taken advantage of its advanced ECL semiconductor technology to build its 1990 SX-3 supercomputer that has four processors and boasts 22 GFLOPS; but the machine is still basically of the von Neumann architecture. In the process area, innovative packaging technologies have been developed to reduce propagation delay of electrical signals and impove on the cooling of the systems.

Most noticeable has been the push by established supercomputer makers for new materials, especially new chips. ECL chips with the latest feature size and the most aggressive packaging are being sort after. Gallium arsenide [3](GaAs) chips, made from a technology that is still in the early stages of the learning curve, is being pursued aggressively by the likes of Cray Research Inc. and Cray Computer Corporation. Cray Research Inc. has actually bought Gigabit Logic, the leading gallium arsenide chipmaker to secure a source for its GaAs chips.

Although fundamentally of the von Neumann architecture, traditional supercomputer products are differentiated. With a few exceptions, each producer started out with software that could only run on its machines and because each wanted to maintain upward compatibility of its own systems and software, evolution of the computers from the fluid state to the specific state produced no standardization and the products remain differentiated.

---

[3]Gallium arsenide (GaAs) is a newer non-silicon chip technology that results in chips that are not only four times faster than their silicon counterparts, but also consume three times less power than chips from the more mature and proven silicon.

Massively Parallel Computers exhibit the characteristics of the Fluid State of the Dynamic Model of Innovation:

The speed at which a sequential von Neumann computer can execute a program depends on several factors that include the architecture of the computer, the speed of the chips in the computer, and the time it takes electrical signals to travel through the wires that connect the different components of the computer. Because the electrical signals travel at about the speed of light (a constant) and the wires through which they travel have some finite length, the speed of light will always be a limiting factor for single processor computers.

Innovations in massively parallel computers (MPCs) are stimulated by the need to get around the performance limitation imposed on von Neumann architectures by the speed of light, and provide the price-performance that users want -- especially for inherently parallel applications -- using readily available components. Only MPCs can provide some of the ultra-high speeds required for many applications. The kinds of performance some MPCs now give some data retrieval or mortgage-backed securities applications, cannot be met by traditional supercomputers.

Numerous MPC designs exits. They vary in the way the processors communicate with each other, the way the memory is connected to the processors, and granularity (the number of processors put on each job). The Goodear Massively Parallel Processor (MPP) design, for example, uses the *nearest-neighbor* configuration with each of its 16 thousand processors connected only to its four nearest neighbors; hence the name *nearest-neighbor-only*. Such a design is very restrictive but very effective for problems with a certain degree of localness as in image processing applications where a spot on the image tends to be affected only by adjacent spots. Thinking Machines' Connection Machine uses the hypercube four-dimensional topology, so the system can be reconfigured to allow any processor to communicate with any other processor in the computer. Many designs abound, with no one company's design being compatible to any other's.

The materials used in MPCs are limited to readily available ones. The chips are from the proven and well-established CMOS silicon technology. No use is made of the chips with the smallest feature sizes (the smaller the feature size the better the performance). Some of the individual processors are actually RISC microprocessors that can be bought of-the-shelf. The disk drives for some of the MPCs are the same readily available and inexpensive commodity disk drives used in personal computers, with the main difference being that many drives are connected in a reflection of the MPC's parallel architecture

In many cases, the entrepreneurs who constitute the "source" of the technology for each design are from academia where, with the backing of potential users, parallel processing ideas have been worked on and espoused for years; the companies founded, have also tended to be located near these institutions. Thinking Machines in Cambridge, near MIT; MasPar Corp. in Sunnyvale near Stanford University; and NCube Corp. in Beaverton, Oregon, near the Intel Systems group which provided the founders of NCube; are testimony to this fact. In some cases, the entrepreneur has come from a larger organization whose formal organizational control didn't suit the entrepreneurial spirit needed for such radical innovations. This was the case with MasPar's Jeff Kalb who left DEC to pursue his MPC ideas. In all the cases, the plants are small scale and located close to the designers and sources or ideas.

Competitive emphasis has been on price/performance and not just on functional performance as the model stipulates. Also, because of software lock-in, major changes are not as frequent as they would be if MPCs were the kinds of products that would become undifferentiated standard products upon evolving into the specific state. This too is an exception to the model as developed for assembled hardware products alone.

Supercomputers and the Utterback & Kim hypotheses Concerning Discontinuous Change in a Product

Cray Research Inc., and the other traditional supercomputer makers exhibit most of the characteristics of the *invaded unit* of the Utterback and Kim Model. They are established competitors in the supercomputer industry whose products, as we have just seen, are in the specific state. They insist that their machines are better than MPCs, because the former have more applications software available and more programmers who are used to the traditional von Neumann programming methods. Moreover, they add, their machines' architectures lend themselves to more applications than the present more specialized MPCs, and that the overall cost of their traditional machines to customers is lower.

As the threat from MPCs has increased, traditional supercomputer makers have redoubled their investments in an effort to improve their established technology. All of them have moved to some limited parallelism (up to 16 processors in some cases) in their designs; but have not taken the larger leap to MPCs. Both NEC and Cray Research Inc. have been using the latest (smallest feature size) in ECL semiconductors, and newly-developed packaging technologies to get the fastest possible chips. Cray has gone one step further, opting to use gallium arsenide chips -- an unproven technology that is not as well understood as silicon, but which can run four times faster than any silicon chip (even ECL) and consume three times less power than ECL. The risk of failure is high. It has also bought Superteck, a minisupercomputer maker, in an attempt to stem errosion of the low-end segment of its market by the invading MPCs and minisupercomputers. These seem just the sort of defensive moves stipulated in the Utterback and Kim model.

On the other hand MPC makers like Thinking Machines, Active Memory Technology, MasPar, and NCube, exhibit the characteristics hypothesized for the invading unit. They are all new entrants whose products exhibit the properties of the fluid state discussed earlier, and who stress the price/performance that their machines offer for inherently parallel applications. [When IBM enters the market, as is rumored, it will be the first established competitor to introduce MPCs.] The performance offered by these systems in some inherently parallel applications like molecular modelling, cannot be matched by any traditional supercomputer. For the long-run, MPC makers also look at their machines as more than just a substitute for the traditional supercomputers. They are working towards broadening the applications of their supercomputers to other areas like artificial intelligence. At the same time established supercomputer makers don't see them as direct competitors.

Prior to starting their companies, MPC founders had acquired superior knowledge in the technology. Daniel Hillis, a Thinking Machines founder, had designed MPCs while a graduate student at MIT, NCube's Stephen Colley came from Caltech, via Sun Microsytems and Silicon Graphics, where the hypercube configuration had been invented, and MasPar is strongly affiliated with academia from Stanford University.

The Utterback Model on the Dynamics of innovation among multiple productive units and the supercomputer industry:

According to the Utterback model, the rate at which firms enter or exit an industry parallels the product innovation in that industry. Figure 4.1 shows the plot of entries and exits by conventional von Neumann supercomputer makers, as well as the total number of these companies in the market at any one time. The shape of this curve is in accordance with the hypotheses on the dynamics of innovation among multiple productive units developed earlier. The number of entries in the mid-1980s increased, contrary to the predictions of the model, as a result of the introduction of vector processing. Vector processing, although a 1970s innovation, actually helped Japanese supercomputer makers Hitachi, NEC, and Fujitsu enter the market in the eighties by attaching vector processors to their mainframe designs and offering the resulting products as supercomputers. IBM re-entered at about the same time by also attaching a vector processor to its top-of-the-line mainframe. All these new entrants are established computer companies with very strong financial positions and other computer, electronics or telecommunications products that can cross-subsidize their supercomputer operations.

In 1989, CDC divested itself of its ETA Systems group, thus getting out of the traditional supercomputer market. Cray Research split off a separate entity, Cray Computer Corporation and moved it to Colorado Springs, Colorado to keep working on traditional supercomputers.

Figure 4.2 shows the entries and exits for minisupercomputer makers. The minisupercomputer innovation was realized primarily because of innovations in semiconductor technology, in particular, developments in

CMOS chip technology that made available chips with many functions per chip, and that consume less power than the ECL chips used in the conventional sequential supercomputers. The minisupercomputer innovation was followed by a wave of entries. The sharp decline cited by the model appears to have started, signalling the emergence of a dominant design. In 1989, Stellar and Ardent merged their minisupercomputer activities to form a new company called Stardent, while Scientific Computer Systems exited. So far in 1990, Supertek has been acquired by Cray Research Inc. and Multiflow has exited the minisupercomputer market.

Although the concept of massively parallel computers has been around for a while, it was in 1983 that several start-ups, some spurred by DARPA funds for parallel processing, entered the supercomputer market with their innovations. Again, a mature chip industry that provided cheap and readily available chips was instrumental to commercialization of the massively parallel processing concept. As figure 4.3 shows, the wave of entries appears to have peaked and the decline started. In 1989, Evans and Suntherland got out of the massively parallel computer (MPC) market, but is still in its other computer businesses. So far in 1990 (by the end of April), Ametek and Floating Point Systems (FPS) have exited the MPC market; FPS is still strong in the minisupercomputers.

The total number of participants for each of the three supercomputer categories is shown in Figure 4.4. For minisupercomputers and MPCs, all the exits in the early part of 1990 may be the beginning of the expected wave of rapid decline which should follow the emergence of dominant designs. While traditional supercomputers do not show any recent exits, we saw earlier that this category already exhibits the characteristics of a productive unit in the specific state and is being strongly invaded by both minisupercomputers and

MPCs; minisupercomputers from the low end (performance-wise) and MPCs from the high end. We speculate that MPCs will eventually take over not just the high-end segment of the market presently dominated by the traditional von Neumann desgins, but also the low-end that is being invaded by minisupercomputers.

5 Discussion

The data collected, together with the use of the Abernathy & Utterback, Utterback, and Utterback & Kim models, indicate that dominant designs in supercomputers of the massively parallel computer (MPC) architecture are emerging. It is evident that the established von Neumann architecture, even with pipelining, vector processors, and large-grained processors, is not adequate for the demanding speeds of supercomputers. The tremendous computer speed gains registered over the years, primarily as a result of the advances in silicon chips, are no longer enough. The speed at which electricity travels through wires is becoming the limiting factor to how fast a supercomputer can be. As such, a new architecture that eliminates this physical limit, the massively parallel architecture, should emerge as the dominant supercomputer design. Eventually, its own speed will be limited by processor communications problems that increase with the number of processors.

Diffusion of MPCs is going to take more than just the performance superiority of these computers. It will require the availability of software that takes full advantage of the parallelism of MPCs. Other factors like the bandwagon effect; the potential profits to be made by both adaptors and MPC makers; complementary innovations like workstations that allow better interface and visualization of results; and key events like if or when IBM enters the market, will all influence how fast this architecture is adopted.

The diffusion is likely to be in two stages: The first stage is for applications that are inherently parallel and some of whose adopters don't mind writing some of their own parallel programs to take advantage of the massive parallelism. We are already going through such a state. The second and more critical stage involves moving the MPC into mainstream applications like manufacturing, financial operations and transaction processing.

The MPC design that emerges as the dominant designs will be the ones that gain the most wide usage and hopefully, the ones that can be used for the most applications -- with the most functionality. Such designs will be greatly influenced not just by performance and functionality, but by the complementary assets of their sponsors [Teece 1987]. Software for the new design will rank top on the list of complementary assets. Distribution channels will be just as important. The role of chance events in this fluid stage should never be underestimated.

Architectures like Thinking Machines' that are fine-grained, have reconfigurable processor communications capability, and have distributed memory, lend themselves to both scientific and artificial intelligence applications. NCube Corporations' architecture with larger processors, larger memory and floating point accelerator at each node, can perform *transaction processing* better. What both machines have in common is the hypercube configuration. It would appear that some form of the hypercube topology -- probably with reconfigurable nodes, and each node having a larger, more powerful processor, large memory, floating point accelerators, and the best software offering -- would prevail.

Ultimately IBM and Cray Research will have a lot to say about which MPC designs emerge as the dominant designs. With their powerful distribution channels, other complementary assets, and the confidence that most customers have in them, the influence of these two companies will be critical if they enter the MPC market. Who can forget what IBM did when it entered the personal computer business; or what it is doing now in the workstation market? IBM's entry into the MPC market could do several things: It could accelerate the number of exits per year, establish another standard or *the* standard, and expand the market by lending credibility to the market and bringing IBM customers and other skeptical customers into the bandwagon. The authors forecast that the massively parallel computer architecture will prevail. It is only a matter of time and effort.

Bibliography

Abernathy, W. J., and J. M Utterback, "Patterns of innovation in technology." *Technology Review* 80:7(June-July):40-47. 1978.

Anderla, Georges, and Dunning, Anthony, "*Compute Strategies 1990-9: Technologies--Costs--Markets.*"

Dorfman, Nancy S., "*Innovation and Market Structure: Lessons from the Computer and Semiconductor Industries.*" Ballinger 1987

Dowline, Second Quarter 1989.

Els, Sharon A., "Configuration and Competitiveness in the Supercomputer Industry." Master's Thesis at MIT. 1989

Fernbach, Sidney "Supercomputers--Past, Present, Prospects." *Future Generation Computer Systems*, July, 1984.

Foster, Foster. "*Innovation : The attacker's Advantage.*" 1986

Freeman, Christopher, "The Economics of Industrial Innovation." MIT Press. Cambridge 1986. Second Edition.

Hillis, Daniel M. *The Connection Machine.* MIT Press. Cambridge, 1985

Hwang, Kai "Multiprocessor Supercomputers for Scientific/Engineering Applications," *Computer*, June, 1985

Kelley, Lauren, "U.S. Supercomputers expand as International markets grow." *Business America*, April 11, 1988.

McCormack, J. S., "Technological Discontinuities: The emergence of Fiber Optics." S.M Thesis at MIT. 1987.

Mueller, D. C., and Tilton, J. E., "R&D Cost as a Barrier to Entry." *Canadian Journal of Economics* vol. 2, no. 4., November1969

Smith and Alexander, *Fumbling Away the Future* (William Marrow and Companies 1988);

Supercomputing -- An Informal Glossary of Terms, IEEE, 1111 19th Street, NW, Washington, DC, 1987

Teece, David J., "Profiting from Technological Innovation: Implications for Integration, Collaboration, Licencing and Public Policy." *The Competitive Challenge.*, edited by David Teece. Ballinger. Cambridge 1987.

Tushman, Michael L., and Anderson, Phillip, "Technological Discontinuities and Organizational Environments." *Administrative Science quarterly,* 31 (1986): 439-465

Utterback, James M. "Innovation and Industrial Evolution in Manufacturing Industries." in *Technology and Global Industry: Companies and Nations in World Economy* edited by Bruce R. Guile and Harvey Brooks. National Academy Press. Washington D. C. 1987. To be elaborated on in *The Dynamics of Innovation in Industry*: Chapter 5. Harvard Business School Press, 1991

Utterback, James M., and Kim, Linsu "Invasion of a Stable Business by Radical innovation." *The Management of Productivity and Technology in Manufacturing,* by Kleindorfer, P R. Plenum Press 1986.

Utterback, James M., and Nolet, T. C. "Product and Process Change in non-assembled products." MIT nternal working paper, Feb.1987.

von Hippel, Eric, *"The sources of Innovation."* Oxford University Press, 1988

Ware, W., "The Ultimate Computer, *IEEE Spectrum* March 1973. pages 89-91

Interviews with over thirty people from the supercomputer industry.

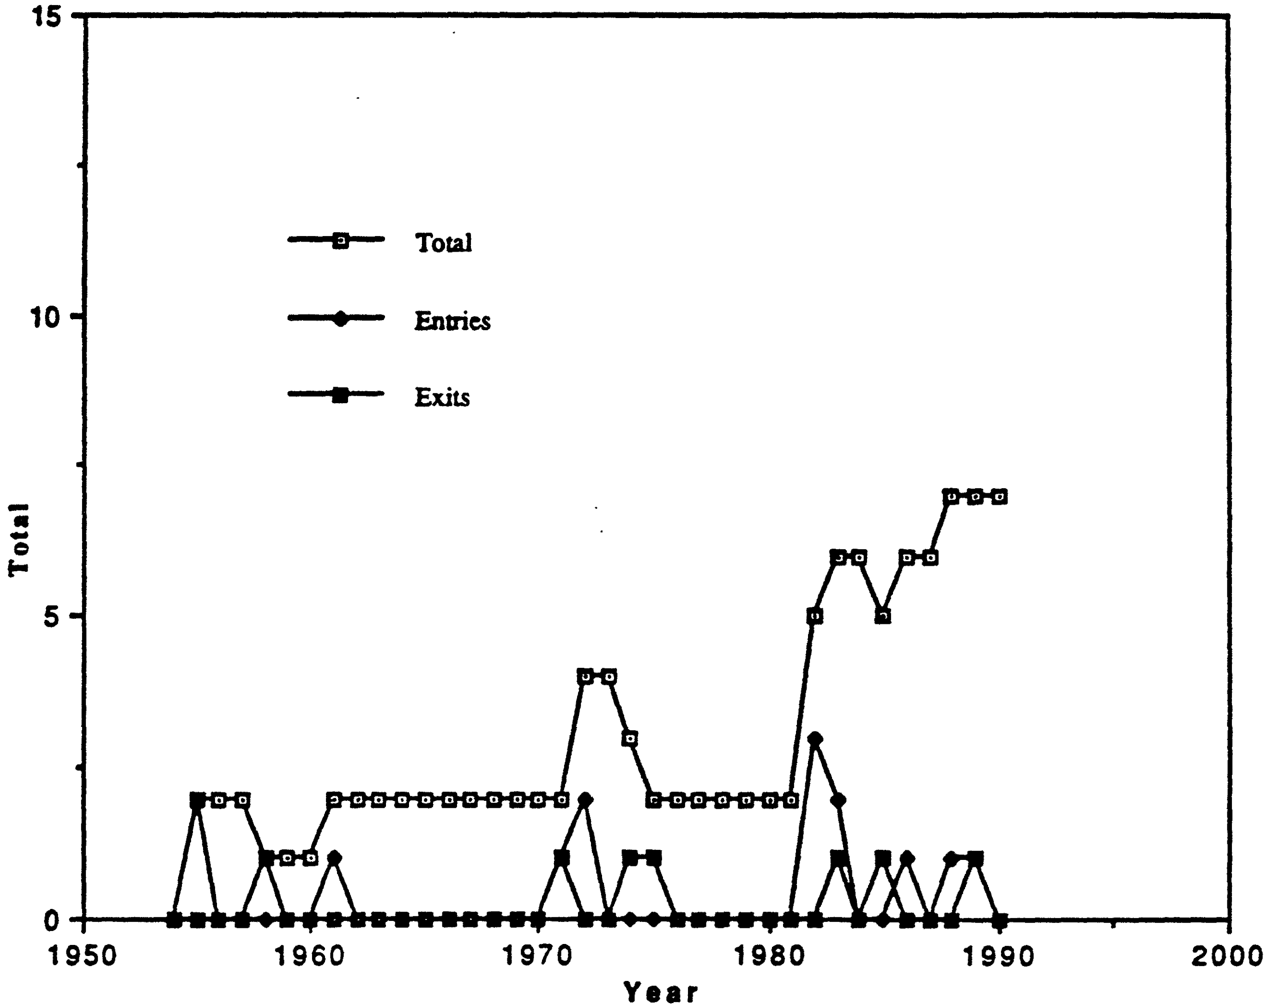

Figure 2.1 : Entry/Exit for Traditional Sequential supercomputers

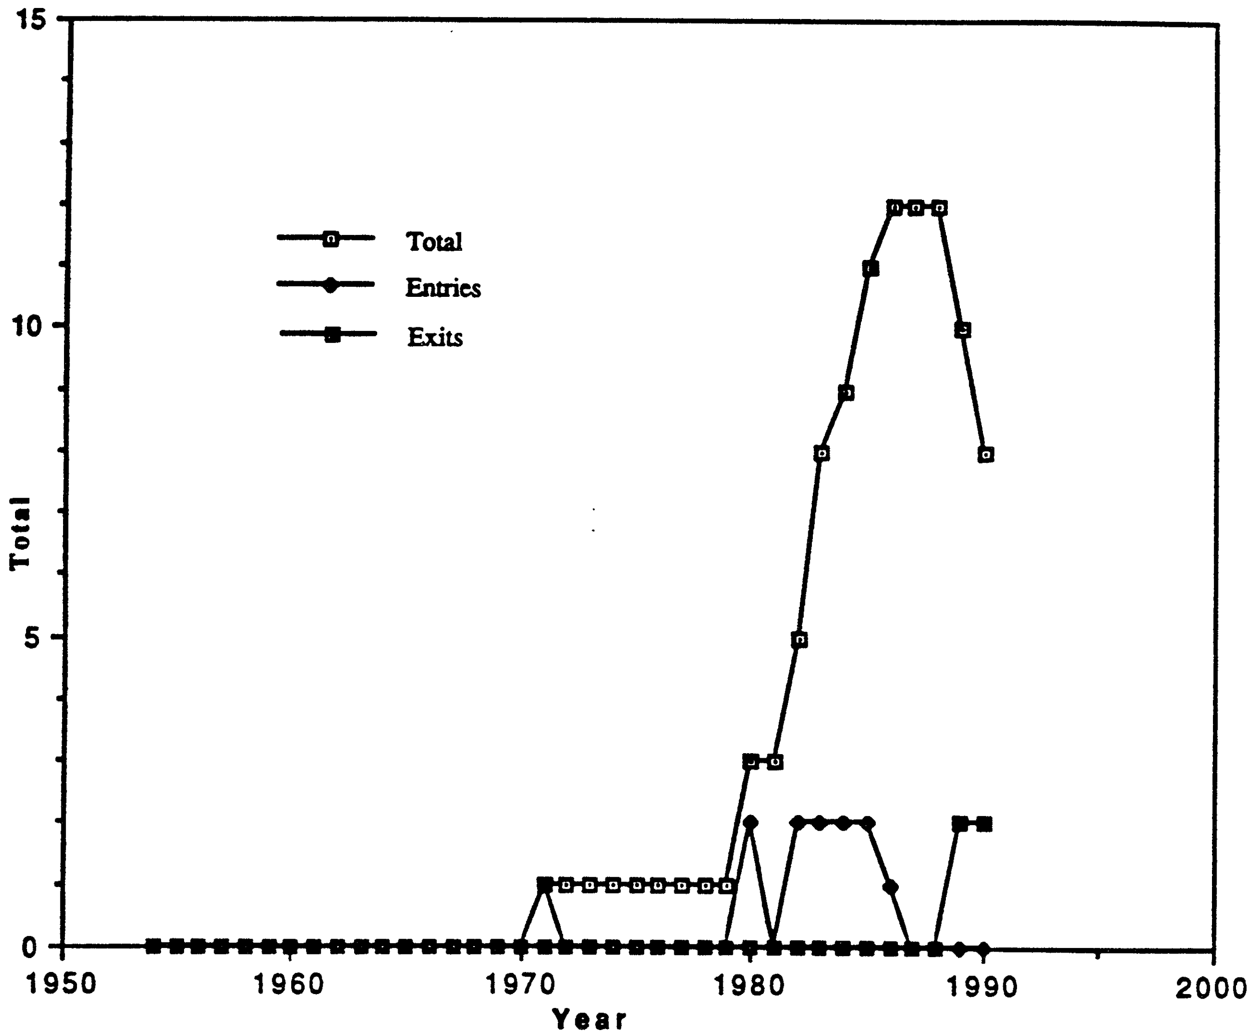

Figure 4.2 : Entry/Exit for Minisupercomputer

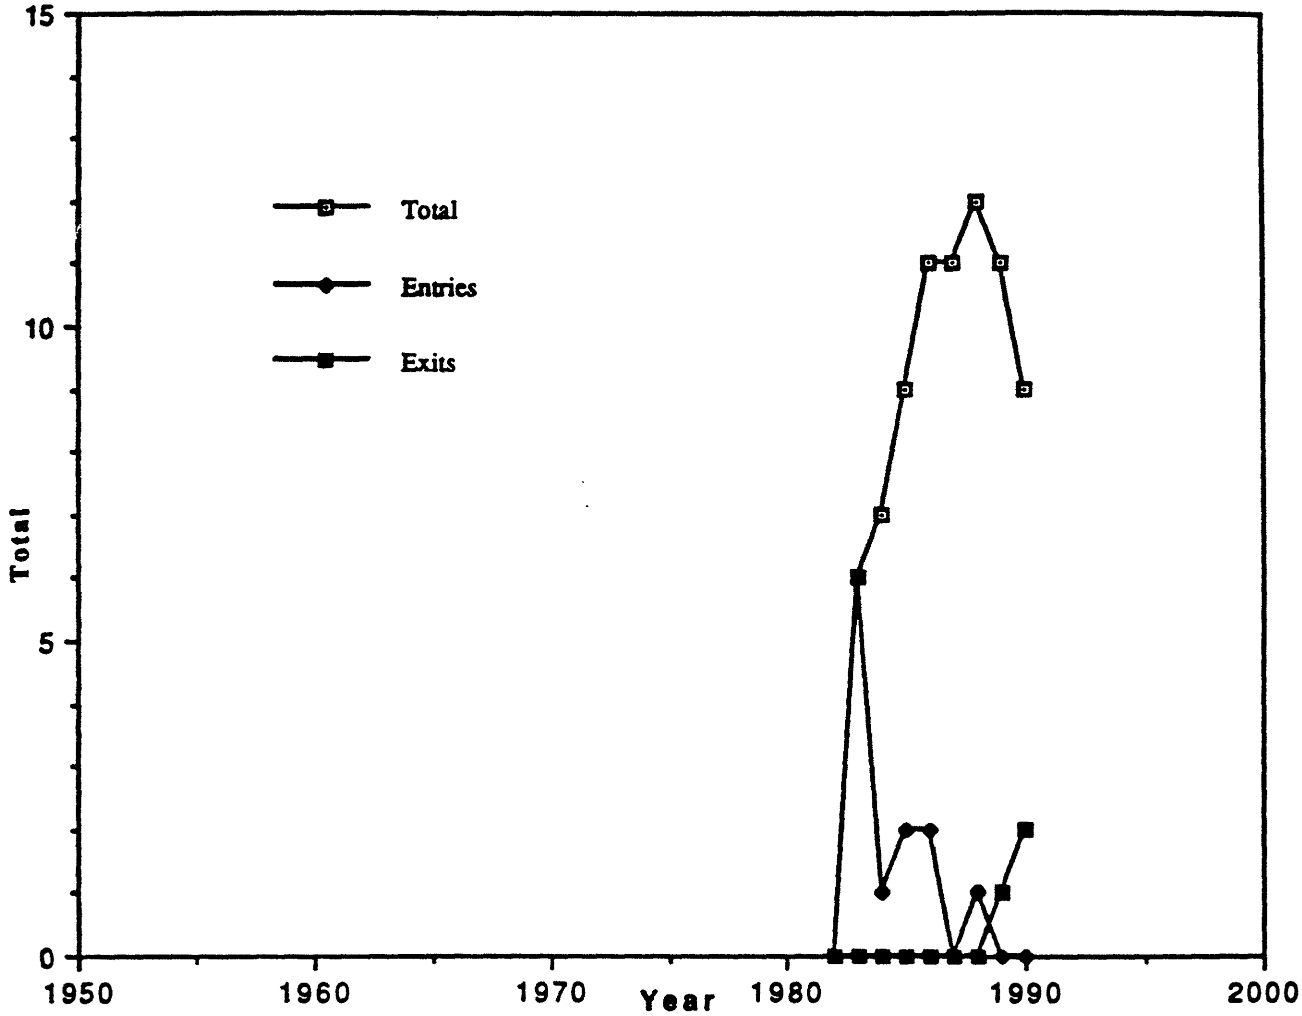

Figure 4.3 : Entry/Exit for Massively Parallel Computers

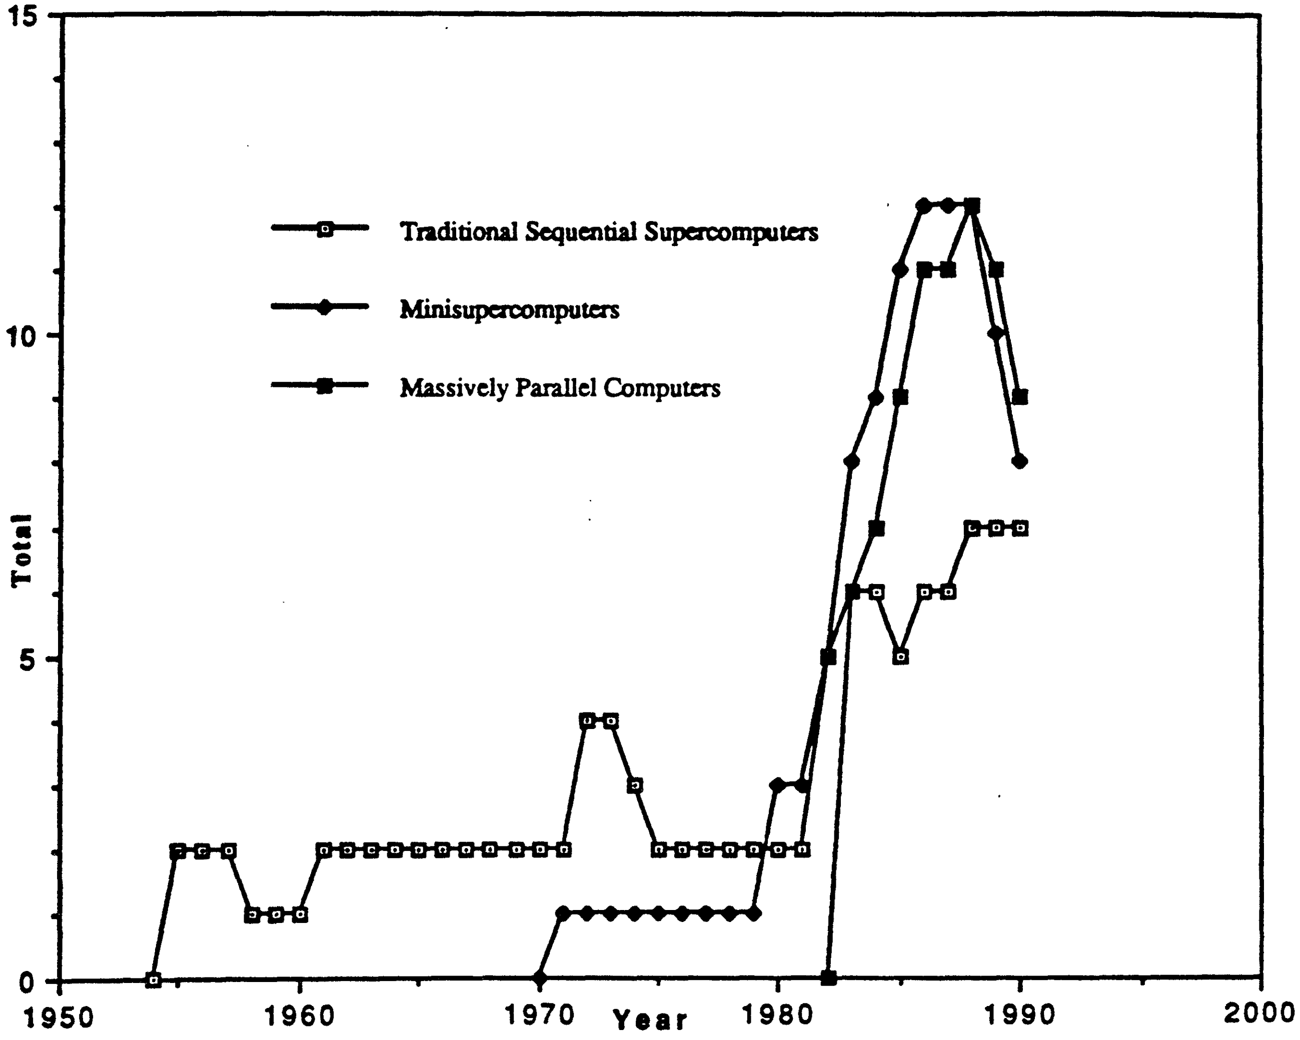

Figure 4.4 : Competing Productive Units